arXiv:1606.06146v1 [q-bio.PE] 20 Jun 2016

Modeling the distribution of distance data in Euclidean space

Ruth Davidson, Joseph Rusinko, Zoe Vernon, and Jing Xi

Abstract. Phylogenetic inference-the derivation of a hypothesis for the common evolutionary history of a group of species- is an active area of research at the intersection of biology, computer science, mathematics, and statistics. One assumes the data contains a phylogenetic signal that will be recovered with varying accuracy due to the quality of the method used, and the quality of the data.

The input for distance-based inference methods is an element of a Euclidean space with coordinates indexed by the pairs of organisms. For several algorithms there exists a subdivision of this space into polyhedral cones such that inputs in the same cone return the same tree topology. The geometry of these cones has been used to analyze the inference algorithms. In this chapter, we model how input data points drawn from DNA sequences are distributed throughout Euclidean space in relation to the space of tree metrics, which in turn can also be described as a collection of polyhedral cones.

1. Introduction to Phylogenetics

A *phylogeny* is a mathematical model of the common evolutionary history of a group of *taxa* X, where X may be a set of genes, species, or multiple individuals sampled within a population. Phylogenies are commonly represented by a *phylogenetic tree* T which is a connected, acyclic graph in which the degree-one vertices, or *leaves* of T, are labeled with the elements of X. The tree T is a representation of the evolutionary history of the set X [**33**].

In many phylogenetic reconstruction pipelines one creates a nucleotide or protein sequence alignment, uses the alignment to estimate pairwise dissimilarities between taxa, and then uses the set of dissimilarities to infer a phylogenetic tree. One class of methods used in the third step of such a pipeline are distance-based algorithms: examples include Neighbor Joining (NJ) [**32**], Unweighted Pair Group Method with Arithmetic Mean (UPGMA) [**35, 36**] (Algorithm 3.1), and FastME [**14**]. These algorithms seek an edge-weighted tree such that the pairwise distances between taxa on the tree reflect the dissimilarities computed from the data. While not as popular as maximum-likelihood methods, distance-based algorithms remain relevant for reasons beyond their speed relative to other methods. For example,

2010 *Mathematics Subject Classification.* 05B25, 05C05, 46N30, 62-07, 62P10, 65C60, 92B10, 97K80.

they serve as components of other popular methods, such as the species tree inference method NJst [**27**] and the DNA sequence alignment method MUSCLE [**16**].

It is possible to analyze the theoretical properties of distance-based algorithms by considering which sets of input data return particular output trees. The set of input data can be viewed as a subset of Euclidean space. It was shown geometrically that NJ solves its optimization problem well in low dimension in [**17**], and in [**11, 12**], geometry provided new insights into biases in the structure of trees reconstructed from phylogenetic data [**2**]. In this chapter we use computational tools in a geometric setting to observe that pairwise dissimilarity data computed from real sequence data will not be uniformly distributed across Euclidean space, but instead will be clustered around an object corresponding to pairwise distances between taxa induced by phylogenetic trees. We derive models for the distribution of this data and discuss possible applications for these models and tools.

2. The Space of additive distance matrices embedded in $\mathbb{R}^{\binom{n}{2}}$

2.1. Dissimilarity maps used in distance-based phylogenetic inference as points in Euclidean space.

DEFINITION 2.1. For a fixed set X of n taxa, a *dissimilarity map* is an $n \times n$ symmetric matrix δ where $\delta(u, v)$ is a measure of evolutionary difference between taxa u and v.

There are many popular ways to estimate such a measure that assume models of sequence evolution such as Jukes-Cantor (JC) [**24**], Kimura two-parameter (K2P) [**25**], or the general time-reversible model (GTR) [**22**]. For example, the JC model of DNA sequence evolution corresponds to a measure of dissimilarity between two DNA sequences of the same length representing taxa u and v:

$$\delta(u, v) = -\frac{3}{4} \log \left(1 - \frac{4}{3}\left(\frac{h}{l}\right)\right) \tag{2.1}$$

where h is the number of places that the sequences differ and l is the shared length of the two sequences. Equation 2.1 is known as the *Jukes-Cantor Correction* because it returns a measure of dissimilarity between taxa based on the expected number of mutations given by the JC model of sequence evolution instead of the observed number of differences in the sequences.

Dissimilarity maps are the inputs to distance-based algorithms. By our definition, $\delta(u, v) = \delta(v, u)$ for all pairs $\{u, v\} \subset X$. Furthermore, $\delta(u, u) = 0$ because δ is a measure of evolutionary distance. So each dissimilarity map δ is also a vector in the Euclidean space $\mathbb{R}^{\binom{n}{2}}$ with coordinates labeled by all distinct pairs of taxa. Finally, we can assume δ is in the positive orthant, which we denote $\mathbb{R}^{\binom{n}{2}}_{\geq 0}$; as we see in Equation 2.1, popular dissimilarity measures such as JC, K2P, and GTR return positive numbers as estimates of evolutionary difference.

The output of a distance-based algorithm is an additive distance matrix (sometimes called a tree metric, though this term has more than one usage in the literature). An *additive distance matrix*, which we will denote as $d(u, v)$ or simply d, can be realized by a phylogenetic tree T with edge weights that are real numbers, where $d(u, v)$ for two distinct elements u and v of X is the sum of the weights of the edges in the unique path between u and v in T. *Ultrametric* additive distance

matrices can be realized by a tree with a root vertex r that is not labeled with a element of X, and satisfy the property that for every two elements u and v of X, $d(u, r) = d(v, r)$. Figure 1 shows an ultrametric on three taxa.

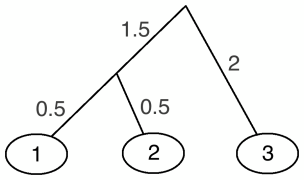

FIGURE 1. An ultrametric on three taxa

2.2. Polyhedral descriptions of the outputs of distance-based methods as points in Euclidean space. Recall that a *polyhedral cone* is the nonnegative span of a set of vectors in $\mathbb{R}^d$ for $d \geq 1$, commonly referred to as the $\mathcal{V}$-representation. Since polyhedral cones are polyhedra, they also have an $\mathcal{H}$-representation as the closed (in the traditional sense of the usual topology on $\mathbb{R}^d$) intersection of a finite number of half-spaces defined by hyperplanes. A *face* of a polyhedral cone is the intersection of the cone with a hyperplane that defines a half-space in $\mathbb{R}^d$ that entirely contains the cone. See [**42**] for a comprehensive introduction to polyhedral theory.

DEFINITION 2.2. A *fan* is a family $\mathcal{F}$ of polyhedral cones in $\mathbb{R}^d$ such that:

(1) if $P \in \mathcal{F}$ then every nonempty face of P is in $\mathcal{F}$, and
(2) if $P_1, P_2 \in \mathcal{F}$ then $P_1 \cap P_2$ is a face of both P_1 and P_2.

Informally, polyhedral fans are special collections of cones that are easy to work with because they "play well together." The set of all additive distance matrices on n taxa, denoted $\mathcal{T}_n$, and the set of ultrametric distance matrices on n taxa, denoted $\mathcal{ET}_n$, are each classified as both a polyhedral fan and a tropical variety. These spaces are studied in [**3**] and [**37**], respectively. We follow the notation of [**12**] to describe these spaces and note that $\mathcal{ET}_n \subset \mathcal{T}_n \subset \mathbb{R}^{\binom{n}{2}}$ for all $n \geq 3$. In other words, the spaces of additive matrices $\mathcal{ET}_n$ and $\mathcal{T}_n$ are *embedded* in the Euclidean space of dissimilarity maps in a very natural way.

In this chapter we model how three-taxon samples of distance data drawn from dissimilarity maps on n taxa are distributed in relation to $\mathcal{ET}_3$. We model this data in relation to $\mathcal{ET}_3$ rather than $\mathcal{T}_3$ because (1) most of the restrictions of the samples to three taxa already lie in $\mathcal{T}_3$ and (2) the condition of ultrametricity is of particular interest to biologists as it indicates that the molecular clock assumption is valid on the sample, which means that evolution occurred at a constant rate over time throughout the entire phylogeny.

We write the restriction of a dissimilarity map to a three-taxon sample δ as a point $(x, y, z) = (\delta(u, v), \delta(u, w), \delta(v, w))$ in $\mathbb{R}^{\binom{3}{2}} = \mathbb{R}^3$. We wish to understand geometric properties of the distribution of points δ in relation to $\mathcal{ET}_3$.

The space $\mathcal{ET}_3$ (see Figure 2) is a two-dimensional polyhedral fan consisting of three two-dimensional cones

$$\text{Span}_{\geq 0}\{(1, 1, 1), (0, 1, 1)\},\ \text{Span}_{\geq 0}\{(1, 1, 1), (1, 0, 1)\},$$

and

$$\text{Span}_{\geq 0}\{(1,1,1),(1,1,0)\},$$

whose pairwise intersections consist solely of the one-dimensional cone $\text{Span}_{\geq 0}\{(1,1,1)\}$. We refer to the two-dimensional (top-dimensional) cones in the fan $\mathcal{ET}_3$ as the *wings*, where

$$\text{Wing } 1 = \text{Span}_{\geq 0}\{(1,1,1),(0,1,1)\},$$
$$\text{Wing } 2 = \text{Span}_{\geq 0}\{(1,1,1),(1,0,1)\},$$

and

$$\text{Wing } 3 = \text{Span}_{\geq 0}\{(1,1,1),(1,1,0)\}.$$

We refer to the one-dimensional cone in $\mathcal{ET}_3$ $\text{Span}_{\geq 0}\{(1,1,1)\}$ as the *spindle*.

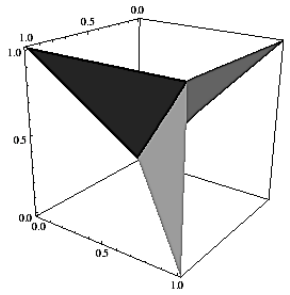

FIGURE 2. Ultrametric outputs $\mathcal{ET}_3$

We introduce a new set of coordinates which describe a dissimilarity map in relation to the geometry of $\mathcal{ET}_3$. Let δ be a dissimilarity map. The relationship of δ to $\mathcal{ET}_3$ can be described by (1) W, the nearest wing to δ, (2) D_1, the signed distance from δ to the nearest wing, (3) D_2, the distance from the point on the nearest wing to the spindle, and (4) D_3, the distance from the point on the spindle to the origin $(0,0,0)$. Then δ can be written as a sum of three pairwise-orthogonal vectors with lengths D_1, D_2, and D_3 representing a path from δ to the origin that begins with a choice of the nearest wing. We call the vector $(W,\ D_1,\ D_2,\ D_3)$ the *path trace* of δ, where $W = \text{Wing } i,\ i \in \{1,2,3\}$. The path trace is an equivalent representation of δ as δ can be uniquely determined by the path trace.

When analyzing a distance-based algorithm, one should emphasize dissimilarity maps likely to arise as outputs of a distance-approximation method based on a popular statistical model of sequence evolution such as JC, K2P, and GTR, which were introduced in Section 1. In the following section we will review distance-based methods and see that they also induce polyhedral structures; studying these structures has shed light on the behavior of the algorithms themselves.

3. Geometry of distance-based methods

3.1. Optimization problems associated to distance-based methods.

PROBLEM 3.1. The Least-Squares Phylogeny problem (LSP) asks, for a given dissimilarity map $\delta \in \mathbb{R}^{\binom{n}{2}}$, what is the additive distance matrix d that minimizes the ordinary Euclidean distance given by the formula

$$\sqrt{\sum_{x,y\in X}(\delta(x,y)-d(x,y))^2}.$$

The distance-based methods Unweighted Pair Group Method with Arithmetic Mean (UPGMA) [**35, 36**] and Neighbor-Joining (NJ) [**32**] are approximations to LSP [**19, 32**]. However, an alternative interpretation of NJ [**21, 15**] is that NJ performs a heuristic search, guided by a linear transformation of δ known as the *Q-criterion* at each agglomeration step, that minimizes a tree-length estimate due to Yves Pauplin [**29**] known as the "Balanced Minimum Evolution" (BME) criterion.

Day showed that both Problem 3.1 and finding the tree minimizing the BME criterion are NP-hard [**13**]. Thus polynomial-time, heuristic, distance-based algorithms such as BIONJ [**20**], Weighbor [**5**], and FastME [**14**] remain essential for solving Problem 3.1 as well as for finding a tree minimizing the BME criterion. These algorithms take dissimilarity maps as inputs, may outperform NJ in terms of topological accuracy under certain conditions, and exhibit superior immunity to reconstruction pathologies well-known to biologists such as long-branch attraction [**4**].

3.2. Polyhedral decompositions of $\mathbb{R}^{\binom{n}{2}}$ induced by distance-based algorithms. Given a distance-based algorithm, there exists a closed-form description of the input space for a phylogenetic reconstruction method due to the natural subdivision of said space induced by the decision criteria of the method. This subdivision consists of regions containing dissimilarity maps that return additive distance matrices realized by the same combinatorial type of tree upon application of the method. For example, the distance-based algorithms NJ and UPGMA induce subdivisions of the input space into families of polyhedral cones defined by the sets of linear inequalities defining the decision steps in these algorithms. To illustrate, we review why there is a polyhedral description of the UPGMA algorithm here.

Algorithm 3.1 UPGMA

- Input: a dissimilarity map $\delta \in \mathbb{R}_{\geq 0}^{n(n-1)/2}$ on $[n]$.
- Output: an ultrametric $d \in \mathbb{R}_{\geq 0}^{n(n-1)/2}$.
  (1) Initialize an unordered set partition $\pi_n = 1|2|\cdots|n$, and set $\delta^n = \delta$.
  (2) For $i = n-1, \ldots, 1$ do
    - From partition $\pi_{i+1} = \lambda_1^{i+1}|\cdots|\lambda_{i+1}^{i+1}$ and distance vector $\delta^{i+1} \in \mathbb{R}_{\geq 0}^{(i+1)i/2}$ choose j, k be so that $\delta^{i+1}(\lambda_j^{i+1}, \lambda_k^{i+1})$ is minimized.
    - Set π_i to be the partition obtained from π_{i+1} by merging λ_j^{i+1} and λ_k^{i+1} and leaving all other parts the same. Let $\lambda_i^i = \lambda_j^{i+1} \cup \lambda_k^{i+1}$.
    - Create new distance $\delta^i \in \mathbb{R}_{\geq 0}^{i(i-1)/2}$ by $\delta^i(\lambda, \lambda') = \delta^{i+1}(\lambda, \lambda')$ if λ, λ' are both parts of π_{i+1} and
$$\delta^i(\lambda, \lambda_i^i) = \frac{|\lambda_j^{i+1}|}{|\lambda_i^i|}\delta^{i+1}(\lambda, \lambda_j^{i+1}) + \frac{|\lambda_k^{i+1}|}{|\lambda_i^i|}\delta^{i+1}(\lambda, \lambda_k^{i+1})$$
      otherwise.
    - For each $u \in \lambda_j^{i+1}$ and $v \in \lambda_k^{i+1}$, set $d(u, v) = \delta^{i+1}(\lambda_j^{i+1}, \lambda_k^{i+1})$.
  (3) Return: A combinatorial rooted tree T with edge weights $w : E(T) \to \mathbb{R}_{\geq 0}$ and the ultrametric $d_{T,w}$ realized by T.

The inequalities coming from the decision steps in UPGMA (Algorithm 3.1) are linear combinations of the original input coordinates, and define a family of polyhedral cones that completely partition the input space. This representation of a polyhedral cone in terms of a set of inequalities is known as an *H-representation*. The *V-representation* of these cones-i.e. a description in terms of their extreme rays-is given in [**11**]. Knowledge of both the H- and V-representations of the cones into which UPGMA divides the input space gives a complete discrete geometric description of the input space, allowing for computations and further conjectures about how this discrete geometric description might affect the performance and behaviors of the algorithm in practice.

While UPGMA outputs ultrametrics, and so maps $\mathbb{R}_{\geq 0}^{\binom{n}{2}}$ to cones in $\mathcal{ET}_n$ indexed by combinatorial rooted trees, Neighbor Joining maps $\mathbb{R}_{\geq 0}^{\binom{n}{2}}$ to cones in $\mathcal{T}_n$ indexed by unrooted tree shapes. As is the case with UPGMA, NJ partitions the input space into a family of polyhedral cones with an H-representation arising naturally from the decision criterion in Step (2) of Algorithm 3.1 which are linear inequalities in terms of the coordinates of the input space. Though the problem has been studied in [**18**], a complete description of the V-representation, or extreme rays, for the cones in the polyhedral partition of the input space induced by NJ is still unknown.

3.3. Using polyhedral geometry to analyze distance-based algorithms. The intersection of the polyhedral decomposition of the input space cones with the piece of the hypersphere $\mathbb{S}_{\binom{n}{2}-1}$ in the section of a Euclidean space with all non-negative coordinates, which we denoted as $HS_{\geq 0}$ in Section 2, provides a finite measure of how frequently the algorithm returns a particular topology sometimes called the *spherical volume* of the cone. Estimation of spherical volumes of polyhedral cones has been used to study the behavior of distance-based phylogenetic methods on small numbers of taxa. For example, in [**17**] estimated spherical volumes were used to assess the agreement of regions in the input space for NJ and the Balanced Minimum Evolution (BME) criterion, which as mentioned above is one of the optimization problems NJ seeks to solve. Also, in [**11**], estimated spherical volumes of cones in the input space for UPGMA indicated that UPGMA may be biased against unbalanced tree topologies, and in [**12**] the notion of spherical volume of small neighborhoods of the input space was further used to investigate the behavior of NJ and UPGMA. These findings assume that observed dissimilarity maps are uniformly distributed when projected on the sphere.

3.4. Limitation to the uniform distribution model. A key biological motivation for properly modeling the distribution of distance data is the so-called *rogue taxa phenomenon* in which the inclusion of a taxon in the estimation of a phylogeny results in reduced accuracy. See [**1**] for an example of a recent method designed to deal with this problem. In [**10**] polyhedral geometry was used to investigate the impact of including an additional taxon on the topology of a BME tree. In a simulation the authors showed that by including distances to an extra taxon and therefore lifting the problem to a higher-dimensional space, one could completely transform the BME topology for the set of taxa corresponding to the original distance matrix in a large number of cases, and that the effect worsened as the number of taxa grew.

But in [**41**] the effect of adding one additional taxon to a small tree was studied for a viral data set as a biological analog of the simulation study done in [**10**], which indicates that this effect may have been overstated due to the model of distances used in the simulation in [**10**]. Additional study of this effect using biological data may lead to more speculation about the true frequency of this effect. Motivated by the failure of the uniform distribution to capture the observed biological features we introduce two families of distribution functions in Section 4 that may better model the dissimilarity maps computed from DNA sequence data.

4. Modeling dissimilarity maps using statistical distributions and geometry

4.1. Two model families for dissimilarity maps. We investigate two families of models of dissimilarity maps based on the underlying structure of $\mathcal{ET}_n$. To construct our first model we use biological data as input, and fit a distribution function that describes the individual coordinates of the path trace associated to the dissimilarity map. We refer to this viewpoint as Model Family (1). In the second model of distance data we assume the data has been generated under the Yule-Harding model with noise accounting for the non-ultrametric tree-like features in the data. Precisely, we use a normal random variable to account for the regular Euclidean distance of an input point from the space $\mathcal{ET}_n$. We then compute the induced distribution of the associated coordinates of the path trace. We refer to this as Model Family (2).

4.2. Using biological data to develop the two model families. To ensure the comparison of these two models reflects biological data we fit the models to data drawn from the biological literature. TreeBASE [**23**] is a public, open-access website maintained by the Phyloinformatics Research Foundation, Inc. Data publicly available from TreeBASE is restricted to data associated to publications that have been submitted for peer review. TreeBASE includes gene, species, and population phylogenies as well as the data used to infer these phylogenies. We used data matrices representing sequence alignments of nucleotide data, inferred distances from these data matrices using the open-source software package MEGA 5 [**40**], and then extracted triples of distances from the large distance matrices to create large samples of dissimilarity maps in three dimensions.

In this chapter we highlight the findings for two data sets we obtained from TreeBASE. In [**28**] the authors are testing support for the classification of Boletineae as a sub-order as well as the families Boletaceae and Paxillaceae, and they find support for these using a phylogenetic analysis of three different genes. The paper [**38**] uses 9 genes to study the phylogenetic relationships of 29 species all within the same family: 22 of the taxa are in the *Daucus* genus and 7 are from related genera. All 29 species belong to the Umbelliferae family.

5. Developing Model Family (1)

It is equally likely that a data point δ is closest to any of the three Wings, so $p_W(\text{Wing } i) = \Pr(W = \text{Wing } i) = 1/3$, $i \in \{1, 2, 3\}$. We model D_1 using a truncated normal distribution, and D_2 and D_3 are modeled in Model (1)-(A) using truncated generalized extreme value distributions (EVDs). For a comprehensive introduction to extreme value theory, see [**9**]. For Model (1)-(A), we designate the

probability density function f_{D_3} for D_3 to be the truncated EVD with parameters α_s and β_s on domain $(0, +\infty)$ (the truncation forbids negative values):

$$f_{D_3}(d_3|\alpha_s, \beta_s) = \frac{e^{\frac{\alpha_s - d_3}{\beta_s}} - e^{\frac{\alpha_s - d_3}{\beta_s}}}{\beta_s c_{\alpha_s, \beta_s}} I_{\{x>0\}},$$

where c_{α_s,β_s} is the normalizing constant and I is the indicator function.

Next, $f_{D_2|D_3}(d_2|d_3, \theta)$ follows the resulting distribution obtained by assuming that $r_1 = \frac{d_2}{d_3}$ follows an EVD with α and β on domain $(0, 1]$, where the truncation is determined by the geometric setting. Note that, conditional on $D_3 = d_3$, $f_{D_2|D_3}$ also follows a truncated EVD. We consider r_1 first and let $d_2 = r_1 d_3$ so that the estimation of α and β does not depend on the actual value of d_3. Similarly we define $r_2 = \frac{d_1}{\min(d_3 + d_2/\sqrt{2},\ d_2/2)}$ and model r_2 by the truncated normal distribution with mean 0 and variance σ^2 on domain $[-\frac{1}{\sqrt{3}}, \frac{1}{\sqrt{3}}]$. Then $d_1 = r_2 \min(d_3 + d_2/\sqrt{2},\ d_2/2)$ will also follow a truncated normal distribution with independently estimated variance σ^2.

The EVD distributions in Model (1)-(A) are replaced with Gamma distributions in Model (1)-(B). We chose the distributions for Model (1)-(A) and Model (1)-(B) for our model by fitting the path trace to histograms taken from triples of distances estimated from biological data using MEGA [**40**].

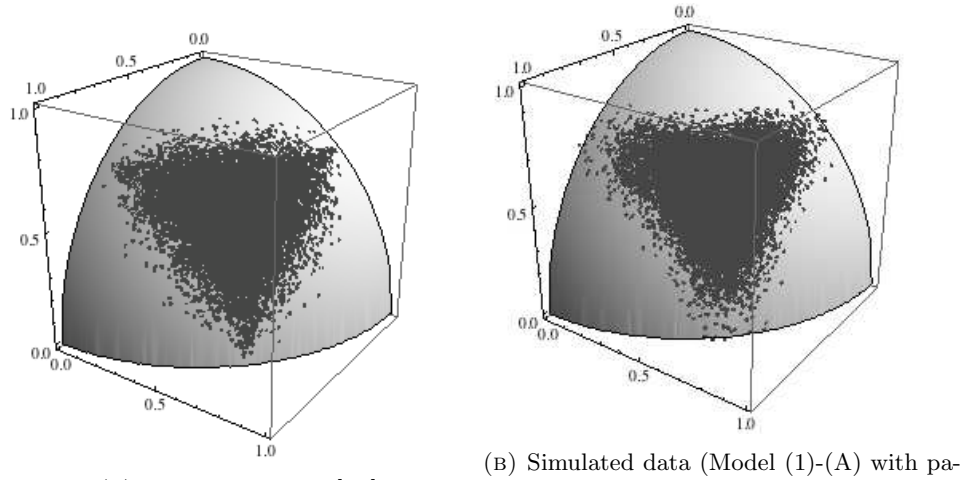

(A) Data points from [**28**]

(B) Simulated data (Model (1)-(A) with parameters estimated from [**28**]

FIGURE 3. A comparison of real versus simulated dissimilarity maps. Data points projected onto the unit sphere.

6. Model Family (2)

To describe Model Family (2), we first describe the Yule-Harding pure birth process (Algorithm 6.1) which generates random rooted, binary phylogenetic trees with leaf set $[n]$. This is also commonly known as the Yule Process but we do emphasize that extinction is often incorporated into the model, in which case it is called a "birth-death" process. We ignore the "death" part of the model for

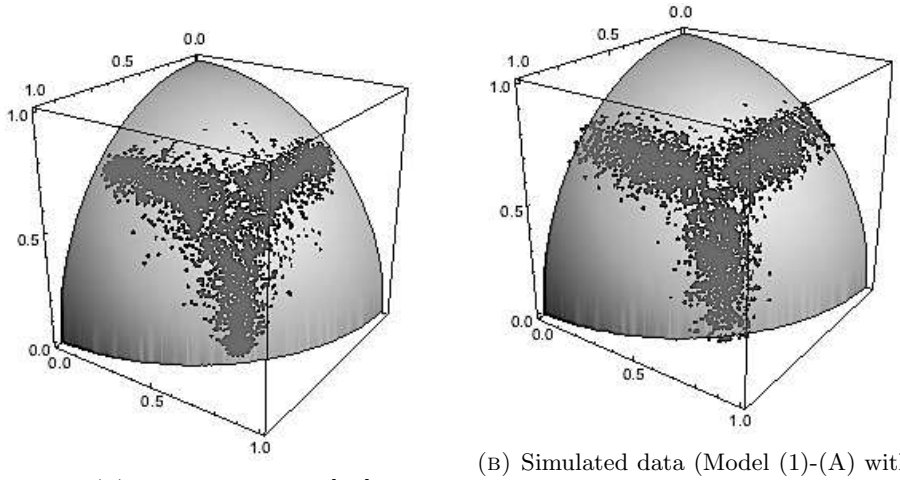

(A) Data points from [**38**]

(B) Simulated data (Model (1)-(A) with parameters estimated from [**38**]

FIGURE 4. A comparison of real versus simulated dissimilarity maps. Data points projected onto the unit sphere.

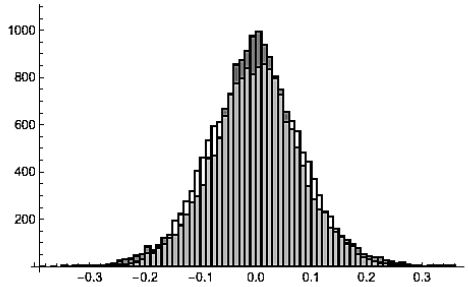

FIGURE 5. Histograms of Real and Simulated Data from Model (1)-(A) $D_1|D_2, D_3$ from [**28**]

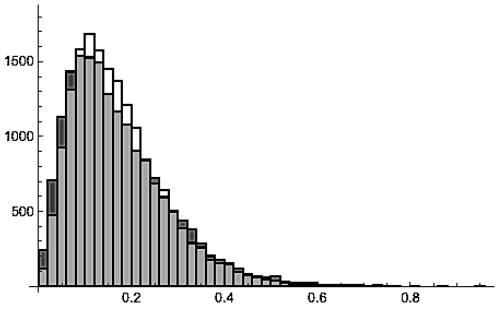

FIGURE 6. Histograms of Real and Simulated Data from Model (1)-(A) $D_2|D_3$ from [**28**]

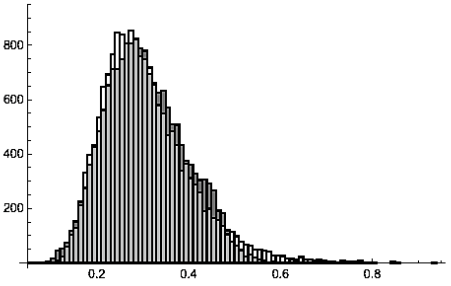

FIGURE 7. Histograms of Real and Simulated Data from Model (1)-(A) D_3 from [**28**]

simplicity. Recall that a *pendant edge* is an edge in a phylogenetic tree incident to a leaf.

Algorithm 6.1 Yule Process

- Input: Leaf set $[n] = \{1, \ldots, n\}$, $n \geq 2$.
- Output: a rooted, binary phylogenetic $[n]$-tree.
- Initialize: Randomly select two leaves x and y from $[n]$ with uniform probability. Identify these two leaves as the leaf set of a rooted binary tree T_1. Set $S_1 = [n] \setminus \{x, y\}$ and $L_1 = \{x, y\}$.
- While $S_i \neq \emptyset$:
  - Randomly select an element x_i of S_i with uniform probability.
  - Randomly select a pendant edge $e_i = (u_i, y_i)$ of T_i with probability determined by the uniform distribution on the set of pendant edges of T_i. Here u_i is a binary internal vertex and $y_i \in L_i$.
  - Subdivide e_i by adding a new vertex v_i.
  - Update T_{i+1} as the tree with leaf set $L_{i+1} = L_i \cup \{x_i\}$, $V(T_{i+1}) = V(T_i) \cup \{v_i, x_i\}$ and $E(T_{i+1}) = E(T_i) \setminus \{e_i\} \cup \{(u_i, v_i), (v_i, y_i), (v_i, x_i)\}$. Update $S_{i+1} = S_i \setminus \{x_i\}$.
- Return T_{n-2+1}, a rooted, binary phylogenetic tree with leaf set $[n]$.

Given a tree T with n leaves generated under the Yule-Harding pure birth model we investigate the distributions for D_2 and D_3 for a subtree obtained by restricting T to three randomly selected taxa. We define D_1 in the same way as for Model Family (1). We determine in this section that D_2 and D_3 should follow Gamma distributions. Throughout this section we will refer to the *graph* distance in a tree T, which is the number of edges in the unique path between a pair of taxa in T.

As mentioned above, dissimilarity maps estimated from DNA sequences is almost never ultrametric, but we use ultrametric trees - i.e. the geometric object $\mathcal{ET}_3$ - as we did for Model Family (1) as a reference for constructing Model Family (2). We again consider the perpendicular projection onto $\mathcal{ET}_3$ of a dissimilarity map δ, so that $\mathcal{ET}_3$ is a baseline object for constructing our model. We consider three cases, which are represented by three labeled tree topologies pictured in Figure 8. We may again, as in Model Family (1) assume that it is equally likely that a data point δ is closest to one of the three wings.

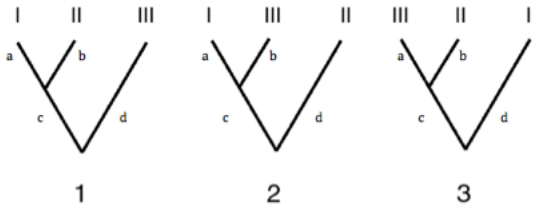

FIGURE 8. Three taxa tree topologies with edge lengths.

The three topologies are illustrated in Figure 8 with branch lengths labeled, where a and b separate the two most closely related taxa, c separates the single internal node of the tree and the root, and d goes from the root to the third taxa.

$$\begin{array}{rcl} \text{Wing } 1 & = & (b+c+d, \frac{2a+b+c+d}{2}, \frac{2a+b+c+d}{2}) \\ \text{Wing } 2 & = & (\frac{a+2b+c+d}{2}, a+c+d, \frac{a+2b+c+d}{2}) \\ \text{Wing } 3 & = & (\frac{a+b+2c+2d}{2}, \frac{a+b+2c+2d}{2}, a+b) \end{array}$$

Under the assumption that T is ultrametric we can substitute $A = a = b$, $B = c = d - a$. Since these trees are subtrees of an n-taxon tree it is important to note that A and B are likely the sum of multiple branch lengths in the larger tree T. We can use the projections and corresponding ultrametric coordinates to give A and B as functions of the input data.

$$\begin{array}{lll} \text{Wing } 1: & A = \frac{b+c+d}{2} & B = \frac{2a-b-c-d}{4} \\ \text{Wing } 2: & A = \frac{a+c+d}{2} & B = \frac{2b-a-c-d}{4} \\ \text{Wing } 3: & A = \frac{a+b}{2} & B = \frac{2c+2d-a-b}{4} \end{array}$$

6.1. Spindle-origin distance.

PROPOSITION 6.1. Under the Yule-Harding model, we may describe the distribution of the set of unscaled distances of the projection from wing to spindle, D_2, as follows: $D_2|\alpha_{D_2}$ has a Gamma distribution with parameters $\frac{2}{\sqrt{3}}\lambda$ and α_{D_2}, where λ is the birth rate and α_{D_2} has expected value $E[\alpha_{D_2}] = 3(\mu_n - 2(1 - \frac{\mu_n}{n-1}))$ where $\mu_n = \sum_{j=2}^{n} \frac{1}{j} \approx ln(n) + \gamma + \frac{1}{2n} - \frac{1}{12n^2}$, and $\gamma \approx 0.5772156649$ denotes the Euler-Mascheroni constant.

PROOF. The sum of n independent exponential random variables with parameter λ is a Gamma distribution with parameters $\alpha = n$ and λ [**26**]. We use this to show $D_2|\alpha_{D_2}$ follows a Gamma distribution, and then determine the expected value of α_{D_2} based on the graph distances between the three taxa.

Suppose we fix a three-taxon subtree T_0. We compute D_2 in terms of the branch lengths of T_0 by finding the Euclidean norm of the orthogonal projection from the point on the wing to the spindle. We find that the D_2 is given by

$$\frac{2}{\sqrt{3}}D,$$

where $D = 2A + 2(A + B) = a + b + c + d$. We note that D is the sum of the branch lengths of T_0. Let the number of branches included in D be α_{D_2}. Under the Yule-Harding model branch lengths follow an exponential distribution with parameter

λ. Thus, D is the scalar multiple of the sum of exponential branch lengths, which implies D has distribution $\Gamma(\lambda, \alpha_{D_2})$ and D_2 follows a Gamma distribution under the Yule-Harding pure birth model. Notice that here α_{D_2} is a constant given the fixed subtree.

Then we see that $D = a + b + c + d = \frac{1}{2}(d(\mathrm{I}, \mathrm{II}) + d(\mathrm{I}, \mathrm{III}) + d(\mathrm{II}, \mathrm{III}))$, which implies $E[D] = \frac{1}{2}(E[d(\mathrm{I}, \mathrm{II})] + E[d(\mathrm{I}, \mathrm{III})] + E[d(\mathrm{II}, \mathrm{III})])$. Since the subtree is randomly chosen, we have $E[D] = \frac{3}{2}E[d(\mathrm{I}, \mathrm{II})]$. Therefore $E[D] = E[E[D|\alpha_{D_2}]] = E[\alpha_{D_2}/\lambda] = E[\alpha_{D_2}]/\lambda = \frac{3}{2}E[d_{ij}]/\lambda$, i.e. $E[\alpha_{D_2}] = \frac{3}{2}E[d_{ij}]$, where $E[d_{ij}]$ is the expected graph distance between a random subset of two taxa in an Yule tree with n. Steel and McKenzie [**39**] showed

$$E[d_{ij}] = 2\left(\mu_n - 2\left(1 - \frac{\mu_n}{n-1}\right)\right).$$

Therefore

$$E[\alpha_{D_2}] = 3\left(\mu_n - 2\left(1 - \frac{\mu_n}{n-1}\right)\right).$$

Additionally we note that the scalar multiple kX of a random variable X that follows a Gamma distribution with parameters α_{D_2} and λ is also a Gamma distribution with parameters α_{D_2} and $k\lambda$. Thus we expect the spindle-to-origin distance to be $\Gamma\left(\alpha_{D_2}, \frac{2}{\sqrt{3}}\lambda\right)$. □

6.2. Wing-spindle distance.

Proposition 6.2. Under the Yule-Harding model, the distribution of the set of distances of the projection from wing to spindle can be described as follows: $D_3|\alpha_{D_3}$ has a Gamma distribution with parameters $2\sqrt{\frac{2}{3}}\lambda$ and α_{D_3}, where λ is the birth rate.

We use similar argument to show $D_3|\alpha D_3$ also follows a Gamma distribution. The Euclidean norm of the projection from the spindle to the origin is $2\sqrt{\frac{2}{3}}B$. For a fixed subtree under the Yule-Harding pure birth model, B is the sum of exponential random variables and therefore follows a Gamma distribution.

In order to determine the expected value for the shape parameter we need to know the expected graph distance in T between the unique internal node and the root in T_0. Steel and McKenzie [**39**] gave the expected value for the distance from the unique internal node of two taxa to the root in T_0. It is possible that a refinement of this work could lead to an explicit formula for the expected value of the shape parameter in terms of the number of taxa. But without such an explicit formula it is simpler to estimate the parameter directly from the data.

7. Data analysis

7.1. Method for scoring the deviation of a model distribution. Given a distribution function for pairwise distance data, one might wish to measure how well it conforms with the data it was designed to model.

In this section we present a method for scoring a model distribution's deviation from a data sample on $HS_{\geq 0}$. We use this method to compare how well our model families, as well as the uniform distribution, model a test data set. We use a numerical test statistic motivated by the assumption that for large samples, the

proportion of the sample in any region of $HS_{\geq 0}$ should be the same as the integral of the proposed density function over that region.

Let S be a sample of N triples of pairwise distance data which has been orthogonally projected onto $HS_{\geq 0}$. Let P be a random sample of triples of distance data drawn from a distribution function γ and then projected onto the unit sphere. We define a test statistic to determine how different S and P are.

We select k randomly generated points v_i on the $HS_{\geq 0}$. Each such point divides $HS_{\geq 0}$, into three spherical triangles

$$T_1 = \{v_i, (1,0,0), (0,1,0)\},$$

$$T_2 = \{v_i, (1,0,0), (0,0,1)\}, \quad \text{and} \quad T_3 = \{v_i, (0,1,0), (0,0,1)\}.$$

For each $1 \leq i \leq k$ we define the triples $\hat{S}$ (resp $\hat{P}$) $= (s_1, s_2, s_3)$ where $s_j = \frac{|S \cap T_j|}{|S|}$. Then we can define the test statistic:

$$\rho(N,k) = \frac{\sum_{i=1}^{k} \sqrt{(s_1-p_1)^2 + (s_2-p_2)^2 + (s_3-p_3)^2}}{k}. \tag{7.1}$$

For large values of N and k, ρ should approach zero.

Note that if a density function f was known for γ, then the p_i could be computed by integrating f along T_1, T_2, and T_3. If f is the density function for the uniform distribution, then as $n \to \infty$, $\hat{P}$ will approach

$$\left(\frac{Area(T_1)}{\frac{\pi}{6}}, \frac{Area(T_2)}{\frac{\pi}{6}}, \frac{Area(T_3)}{\frac{\pi}{6}} \right).$$

7.2. Fitting data sets from Treebase.org to Model Families (1) and (2). We used custom Mathematica software available on our supplementary materials website to test the ability of our different model families to fit biological data sets. The results for ten data sets are shown in Table 1. Rows are indexed by the ten data sets and the entries in the tables are the score based on our geometric test statistic explained in Section 7.1.

We explain the labels for the columns: Models (1)-(A) and (1)-(B) are from Model Family (1) where (1)-(A) indicates the choice of the extreme value distribution and (1)-(B) indicates the choice of the Gamma distribution. Uniform is the model where input points are uniformly distributed in the input space. Lower scores are best, as explained in Section 7.1, and the best score for each data set is in bold text. Rows are indexed by the data matrix name on Treebase.org: M18755 is the data matrix we used from [**28**], and we were no longer able to find the study [**38**] in Treebase.org by the time of this publication, so we have made the distances we calculated from that data matrix available in our online supplementary materials. The row corresponding to that data matrix is labeled "Daucus."

In Table 1, for each data set, we used $k = 10000$ sample points for the area subdivisions in our scoring system. For small data matrices, we used all possible triples of input points $(\delta(u,v), \delta(u,w), \delta(v,w)) = (x,y,z)$, i.e. $N = \binom{n}{3}$ where n was the number of taxa in the matrix. For data matrices with 50 or more taxa, we used approximately $N = 20,000$ data points, where N varied slightly because our method for random sampling of points $(\delta(u,v), \delta(u,w), \delta(v,w))$ produced duplicates.

While the results in Table 1 are interesting, and definitely show that for these datasets, the uniform distribution perform poorly, analysis of many more data sets

would be necessary to establish a clear trend. It also appears that Model (1)-(B) which blends the geometric features of Model Family (1), with the choice of Gamma distribution motivated by Model Family (2), may provide the best overall accuracy.

Data Matrix	Model (1)-(A)	Model (1)-(B)	Model 2	Uniform
M1807	0.0208444	0.0276282	**0.01446**	0.0892374
M1789	**0.00944452**	0.00961268	0.0154542	0.26513
M2566	0.0624382	**0.04826421**	0.0736874	0.060538
M25733	0.0140685	**0.0131326**	0.0334055	0.264626
M18755	0.00826286	**0.00819624**	0.0134024	0.293712
M25665	0.0141961	**0.0141181**	0.0276884	0.237157
Daucus	0.00879329	**0.00672433**	0.0148475	0.294249
M536	0.0251412	**0.019827**	0.0215093	0.130195
M806	**0.0282066**	0.0331788	0.0359551	0.153135
M1169	**0.0115309**	0.0122847	0.0177641	0.328366

TABLE 1. Scores $\rho(N, k)$ from each model for data sets with Jukes-Cantor Distances from Treebase.org. Model of best fit is in bold.

8. Motivations: choice of distributions in Model Families (1) and (2)

The variance of the normal distribution parametrizes D_1 in both model families (1) and (2). A value of $d_1 = 0$ means that a data point is in $\mathcal{ET}_3$, and so is an ultrametric additive distance matrix. Therefore larger d_1 should correspond to higher variations in the rates of evolution between the three taxa for the data point; i.e. deviation from the molecular clock. The normal distribution is thus a natural choice for both model families.

We first chose the EVD to model D_2 and D_3 because histograms of the corresponding steps in path traces of the real data points closely resembled the probability density function for the EVD. Figures 5, 6, and 7 show the histograms of data points of the type $(\delta(u, v), \delta(u, w), \delta(v, w))$ plotted together with histograms of data points generated using Model (1)-(A). However, since both the Gamma and EVD are transformations of the exponential distribution, [**26**], and because of our results from Section 6, we felt it was necessary to include Model (1)-(B).

Our model families seem to be capable of capturing rather subtle features of the biological data beyond how far the data deviate from the molecular clock. We observed that in some datasets, such as in [**38**], data points are less distributed in the center (i.e. around the spindle). This phenomenon is visible in Figures 4a and 4b. This means that the distribution of D_2 has lower density when D_2 is close to 0 and near the boundary. We comment on this feature for this data set because the density functions of the EVD and Gamma distributions can accommodate this feature of the data with appropriate parameter choices in both Model Families (1)-(A) and (1)-(B), while the density function for the exponential distribution on branch lengths used in Model Family (2) cannot.

9. Non-ultrametric trees and application to larger data sets

The model families described below emphasize the geometry of ultrametric trees ($\mathcal{ET}_3$). However, a similar analysis could be conducted relative to the more

general tree space $\mathcal{T}_n$. With only three taxa, every dissimilarity map that satisfies the triangle inequality is already realizable as a pairwise distance on a tree [**34**]. Therefore, it is more interesting to consider the relationship between maps and the more general tree space in the context of $\mathcal{T}_4$.

The geometry of $\mathcal{T}_4$ can be described as the union of three *superwings*

$$W_{14} = \{P \in \mathbb{R}^6 | x_1 2 + x_3 4 = x_{13} + x_{24} \leq x_{14} + x_{23}\},$$

$$W_{13} = \{P \in \mathbb{R}^6 | x_1 2 + x_3 4 = x_{14} + x_{23} \leq x_{13} + x_{24}\},$$

and

$$W_{12} = \{P \in \mathbb{R}^6 | x_1 3 + x_2 4 = x_{14} + x_{23} \leq x_{12} + x_{34}\}.$$

While this model begins with an emphasis on general tree space, we shift our emphasis to the condition of ultrametricity through a projection to $\mathcal{ET}_3$ at a later stage.

In the earlier sections of this chapter we assumed our biological data arose from the addition of Gaussian noise to phylogenetic trees identified as points in Euclidean space. We continue with this paradigm and model the distance of a dissimilarity map to the nearest superwing using a normal distribution. Next we project the image of the point on the superwing to the *superspindle* or set of coordinates with $x_{12} + x_{34} = x_{13} + x_{24} = x_{14} + x_{23}$. This distance could be modeled by a Gamma distribution.

As the sum of branch lengths under the Yule-Harding model we can assume $x_{12} + x_{34}$ follow a Gamma distribution. By knowing this sum we can reconstruct all coordinates using only x_{12}, x_{13}, and x_{23}. These triples of coordinates lie in $\mathbb{R}^3$ and correspond to triples of pairwise distances drawn from a tree on n taxa and can thus be modeled using the findings in Section 4.

Moreover, when there are more than four taxa the geometry of tree space is simply the intersection of superwings over all sets of four taxa. This follows directly from Buneman's four-point condition (or three point condition in the ultrametric setting) [**6**]. While it is possible to build more complicated distributions relating dissimilarity maps to $\mathcal{T}_n$ or $\mathcal{ET}_n$ such models would require a larger set of parameters, and may necessarily contain more information than the minimal cases. However, the following question is paramount for turning the theoretical work of this chapter into a practical tool for computational biologists.

QUESTION 9.1. How does a distribution of dissimilarity maps on subsets of 3 (or 4) of n taxa induce a distribution of dissimilarity maps on n taxa?

10. Applications and future work

The models in this chapter provide new tools for interpreting analytic results about distance-based phylogenetic methods using geometry. Analyses of trends in the shapes of trees using notions such as balance statistics has led to interesting insights and questions about model assumptions in the past [**2**]. The tools presented in Sections 5 and 6 establish a framework for evaluating suitability of model assumptions for phylogenetic data. Recently, the TreeBASE website (accessed June 15, 2016) stated: "as of April 2014, TreeBASE contained data for 4,076 publications written by 8,777 different authors" [**23**]. A systematic analysis of more alignment matrices using the software available in our supplementary materials might give a clearer picture of trends in biological data that relate the model families presented

here, providing empirical support for choosing one model family as more suitable for further development and study.

Our use of the geometric object $\mathcal{ET}_3$ as the backbone of our models serves many purposes. For example, constructing Model Families (1) and (2) around $\mathcal{ET}_n$ provides a unifying framework to study which Model Family is better suited to modeling biological data. The Yule-Harding pure birth process does produce ultrametric trees, which are points in $\mathcal{ET}_n$. So, comparing the Model Families (1) and (2) in this geometric setting may provide insight into the implications of assuming the Yule-Harding model in data simulation.

Furthermore, the Mathematica software in our supplementary materials can be used directly to estimate the deviation of a biological data set from the molecular clock, via the amount of variance in the normal distribution parameter estimated as D_1. We have already called attention to the fact that biological data is rarely ultrametric, but in the biology and computer science communities, especially in the field of phylogenomics, the assumption of a molecular clock is still necessary to provide theoretical guarantees in many instances. For example, in [**30**] it is shown that without the assumption of a molecular clock, no theoretical guarantees bounding gene tree estimation error can be made.

Also, some phylogenomic methods that bypass the need for gene tree estimation rely on the assumption of a molecular clock for their theoretical guarantees. For example, the recently developed method SVDquartets [**7**] relies on the identifiability result in [**8**] for theoretical guarantees, and requires a molecular clock. So, we argue that establishing the deviation of a data set from the molecular clock using our software would be useful in determining whether or not it is reasonable to apply a method that is only guaranteed to work well on data that follows a molecular clock to that specific data set.

The results shown in this chapter are restricted to distances computed under the Jukes-Cantor model of sequence evolution. We performed the same experiments with the K2P model, but the results on the data sets we investigated were so similar for each data set that we omitted these findings. However, future work should include an analysis of both distances computed from amino acid data as well as more general statistical models of sequence evolution such as GTR [**22**] and GTR+Γ [**31**].

10.1. Future Research Questions.

Question 10.1. Is there a correlation between the number of missing taxa in a tree and the parameters of the Gamma distribution in Model Family (2)? How does this shed light on how closely a set of sampled species may be related in the Tree of Life?

Question 10.2. What is the likelihood of the rogue taxa phenomenon using Model Family (2), which more naturally extends to $n > 3$ than Model Family (1)?

Question 10.3. Can one provide an inductive and computationally tractable model for the distribution of distance data for pairwise dissimilarity vectors on n taxa for Model Family (1)? If so, how will this model compare to Model Family (2)

Question 10.4. For sequence data generated under the coalescent model, is there a correlation between the estimated Gamma distribution parameters across the different gene trees? Could the variation among these parameters inform the design of summary methods?

11. Supplementary material

All supplementary materials, including the distance matrixes computed from the alignment data matrices downloaded from Treebase.org, and Mathematica software for generating simulated distance data, can be found at goo.gl/08PUC5.

12. Acknowledgements

R.D. was partially supported by the National Science Foundation (DMS 0954865). J.R. was partially supported by grants from the National Center for Research Resources (5 P20 RR016461) and the National Institute of General Medical Sciences (8 P20 GM103499) from the National Institutes of Health. Z.V. was partially supported by the National Institute of General Medical Sciences (8 P20 GM103499) from the National Institutes of Health. J. X. was partially supported by the David and Lucille Packard Foundation. We also thank James Degnan, Megan Owen, Mike Steel, Seth Sullivant, Caroline Uhler, Tandy Warnow, and Ruriko Yoshida for helpful discussions over the course of this project.

References

1. A. J. Aberer, D. Krompass, and A. Stamatakis, *Pruning rogue taxa improves phylogenetic accuracy: an efficient algorithm and webservice*, Syst. Biol. **62** (2013), no. 1, 162–166.
2. D. Aldous, *Stochastic models and descriptive statistics for phylogenetic trees, from Yule to today*, Stat. Science **16** (2001), 23–34.
3. F. Ardila and C.J. Klivans, *The Bergman complex of a matroid and phylogenetic trees*, J. Combin. Theory, Ser. B **96** (2006), 38 – 49.
4. J. Bergsten, *A review of long-branch attraction*, Cladistics **32** (2005), 163193.
5. W. M. Bruno, N. D. Socci, and A. L. Halpern, *Weighted Neighbor-Joining: a likelihood-based approach to distance- based phylogeny reconstruction*, Mol. Biol. Evol. **17** (2000), 189–197.
6. P. Buneman, *A note on the metric properties of trees*, J. Combin. Theory, S. B **17** (1974), no. 1, 48–50.
7. J. Chifman and L. Kubatko, *Quartet inference from SNP data under the coalescent model*, Bioinform. **30** (2014), no. 23, 3317–3324.
8. ———, *Identifiability of the unrooted species tree topology under the coalescent model with time-reversible substitution processes, site-specific rate variation, and invariable sites*, J. Theoret. Biol. **374** (2015), 35–47.
9. S. Coles, *An introduction to statistical modeling of extreme values*, Springer-Verlag, London, 2001.
10. M.A. Cueto and F. A. Matsen, *Polyhedral geometry of phylogenetic rogue taxa*, Bull. Math. Biol. **73** (2011), 1202–1226.
11. R. Davidson and S. Sullivant, *Polyhedral combinatorics of UPGMA cones*, Advances in Applied Mathematics **50** (2013), 327–338.
12. R. Davidson and S. Sullivant, *Distance-based phylogenetic methods around a polytomy*, IEEE/ACM Trans. Comput. Biol. Bioinf. **11** (2014), no. 1, 325–335.
13. W. Day, *Computational complexity of inferring phylogenies from dissimilarity matrices*, Bull. Math. Biol. **49** (1987), 461–467.
14. R. Desper and D. Gascuel, *Fast and accurate phylogeny reconstruction algorithms based on the minimum-evolution principle*, J. Comput. Biol. **9** (2002), no. 5, 687–705.
15. R. Desper and O. Gascuel, *Theoretical foundation of the balanced minimum evolution method of phylogenetic inference and its relationship to weighted least-squares tree fitting*, Mol. Biol. Evol. **21** (2004), no. 3, 587–598.
16. R. C. Edgar, *MUSCLE: multiple sequence alignment with high accuracy and high throughput*, Nucl. Acids Res. **32** (2004), no. 5, 1792–1797.
17. K. Eickmeyer, P. Huggins, L. Pachter, and R. Yoshida, *On the optimality of the Neighbor-Joining algorithm*, Alg. Mol. Biol. **3** (2008), doi:10.1186/1748–7188–3–5.
18. K. Eickmeyer and R. Yoshida, *Partitioning the sample space on five taxa for the Neighbor-Joining algorithm*, http://arxiv.org/abs/math/0703081.

19. C. Fahey, S. Hosten, N. Krieger, and L. Timpe, *Least squares methods for equidistant tree reconstruction*, (2008), http://arxiv.org/abs/0808.3979.
20. O. Gascuel, *BIONJ: an improved version of the NJ algorithm based on a simple model of sequence data*, Mol. Biol. Evol. **14** (1997), 685–695.
21. O. Gascuel and M. Steel, *Neighbor-Joining revealed*, Mol. Biol. Evol. **23** (2006), no. 11, 1997–2000.
22. X. Gu, Y.X. Fu, and W. H. Li, *Maximum likelihood estimation of the heterogeneity of substitution rate among nucleotide sites*, Mol. Biol. Evol. **12** (1995), 546–557.
23. The Phyloinformatics Research Foundation Incorporated, *TreeBASE web*, 2010, pp. http://treebase.org/treebase–web/home.html.
24. T. H. Jukes and C. R. Cantor, *Evolution of protein molecules*, Mamm. Protein Metab. **3** (1969), 21–132.
25. M. Kimura, *A simple method for estimating evolutionary rates of base substitutions through comparative studies of nucleotide sequences*, J. Mol. Evol. **16** (1980), 111–120.
26. L. M. Leemis and J. T. McQueston, *Univariate distribution relationships*, The American Statistician **62** (2008), no. 1, 45–53.
27. L. Liu and L. Yu, *Estimating species trees from unrooted gene trees*, Syst. Biol. **60** (2011), 661–667.
28. M. E. Nuhn, M. Binder, A. F. S. Taylor, R. E. Halling, and D. S. Hibbett, *Phylogenetic overview of the Boletineae*, Fungal Biol. **117** (2013), 479–511.
29. Y. Pauplin, *Direct calculation of a tree length using a distance matrix*, J. Mol. Evol. **51** (2000), 41–47.
30. S. Roch and T. Warnow, *On the robustness to gene tree estimation error (or lack thereof) of coalescent-based species tree methods*, Syst. Biol. (2015), doi: 10.1093/sysbio/syv016.
31. J. S. Rogers, *Maximum likelihood estimation of phylogenetic trees is consistent when substitution rates vary according to the invariable sites plus Gamma distribution*, Syst. Biol. **50** (2001), 713–722.
32. N. Saitou and M. Nei, *The Neighbor-Joining method: a new method for reconstructing phylogenetic trees*, Molecular Biology and Evolution **4** (1987), 406–425.
33. C. Semple and M. Steel, *Phylogenetics*, Oxford University Press, Oxford, United Kingdom, 2003.
34. C. Semple and S. Steel, *Phylogenetics*, Oxford University Press, Oxford, 2003.
35. R. R. Sokal and C. Michener, *A statistical method for evaluating systematic relationships*, Univ. Kansas Sci. Bull. **38** (1958), 1409–1438.
36. R.R. Sokal and P.H.A. Sneath, *Numerical taxonomy*, W.H. Freeman, San Francisco, 1963.
37. D. Speyer and B. Sturmfels, *The tropical Grassmannian*, Adv. Geom. **4** (2004), 389–411.
38. D. Spooner, P. Rojas, M. Bonierbale, L.A. Mueller, M. Srivastav, D. Senalik, and P. Simon, *Molecular phylogeny of Daucus*, Syst. Bot. **38** (2013), 850–857.
39. M. Steel and A. McKenzie, *Properties of phylogenetic trees generated by Yule-type speciation models*, Math. Biosci. **170** (2001), no. 1, 91–112.
40. K. Tamura, D. Peterson, N. Peterson, G. Stecher, M. Nei, and S. Kumar, *MEGA5: Molecular evolutionary genetics analysis using maximum likelihood, evolutionary distance, and maximum parsimony methods*, Mol. Biol. Evol. **28** (2011), 2731–2739.
41. K.M. Westover, J.P. Rusinko, J. Hoin, and M. Neal, *Rogue taxa phenomenon: a biological companion to simulation analysis*, Mol. Phylogenetics Evol. **1** (2013), 1–3.
42. G. M. Ziegler, *Lectures on polytopes*, vol. 152, Springer Science & Business Media, 1995.

Department of Mathematics, University of Illinois Urbana-Champaign
E-mail address: redavid2@illinois.edu

Department of Mathematics and Computer Science, Hobart and William Smith Colleges
E-mail address: rusinko@hws.edu

Department of Mathematics, Washington University in St. Louis
E-mail address: zoe.vernon@wustl.edu

Department of Mathematics, North Carolina State University
E-mail address: jxi2@ncsu.edu